

# HP LeftHand SAN Solutions

Support Document

---

User Manuals

VSA User Manual for VMware ESX for SANiQ 8.0



## Legal Notices

### Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

### Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

### Copyright Notices

© Copyright 2009 Hewlett-Packard Development Company, L.P.

## Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

**<http://h20230.www2.hp.com/selfsolve/manuals>**

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to:

**<http://h20229.www2.hp.com/passport-registration.html>**

Or click the **New users - please register** link on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

The following table indicates changes made to this document since the last released edition.

### Document Changes

Chapter	Version	Changes

## Support

Visit the HP Software Support web site at:

**[www.hp.com/go/hpsoftwaresupport](http://www.hp.com/go/hpsoftwaresupport)**

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:

**<http://h20229.www2.hp.com/passport-registration.html>**

To find more information about access levels, go to:

**[http://h20230.www2.hp.com/new\\_access\\_levels.jsp](http://h20230.www2.hp.com/new_access_levels.jsp)**


1 Chapter: Designing A Virtual SAN .....	7
Overview .....	7
Getting Started .....	7
Concepts and Terms .....	8
Designing Your Virtual SAN .....	9
Configuration Requirements .....	10
Best Practices .....	10
Hardware Design for VSA .....	12
CPU .....	12
Memory .....	12
Controllers and Hard Drives .....	13
Virtual Network Design .....	14
Using VSAs with LeftHand SANs .....	15
Cloning VSAs .....	16
SAN/iQ Managers on Physical Platforms .....	16
VSA Managers on Physical Platforms .....	17
Performance Impact of Mixing Platforms in Clusters .....	17
Sample Configurations .....	17
Single Node Configuration .....	17
Two-Node Configuration .....	18
Three-Node-Plus Configuration .....	19

This page is intentionally left blank

# 1 Chapter: Designing A Virtual SAN

## Overview

The Virtual SAN Appliance for VMware® ESX Server (VSA™) enables full featured use of the SAN/iQ® software in a virtualized environment. Implementing SAN functionality within server virtualization enables many unique server and storage configurations. When deployed correctly, LeftHand Networks virtual SANs are scalable, highly available, and fully redundant.



**Figure 1** Example of a Virtual SAN using the Virtual SAN Appliance for VMware® ESX Server (VSA)

## Getting Started

Find detailed instructions for properly installing and configuring the VSA in both the VMware VI Client and the SAN/iQ software in the *Quick Start Guide* for the VSA. The VSA *Quick Start Guide* is available in PDF format and is installed in the Documentation subdirectory with the Centralized Management Console program files.

## Concepts and Terms

**Table 1    Glossary for the VSA**

<b>Term</b>	<b>Definition</b>
<b>VMware Terms</b>	
<b>Virtual Switch</b>	VMware Infrastructure lets you create abstracted network devices called virtual switches (vSwitches). A vSwitch can route traffic internally between virtual machines and link to external networks.
<b>Virtual Disk</b>	A file or set of files that appears as a physical disk drive to a guest operating system. These files can be on the host machine or on a remote file system.
<b>VirtualCenter Server</b>	A service that acts as a central administrator for VMware servers connected on a network. This service directs actions on the virtual machines and the virtual machine hosts. VirtualCenter Server is the working core of VirtualCenter.
<b>SAN/iQ Terms</b>	
<b>Cluster</b>	A cluster is a grouping of storage nodes that create the storage pool from which you create volumes.
<b>Failover Manager</b>	A specialized manager running as a VMware guest operating system that allows you to place a quorum tie-breaker node into a 3rd location in the network to provide for automated failover/failback of the Multi-Site SAN clusters. The Failover Manager is designed to run on VMware Server, which is installed on your own hardware, separate from the SAN hardware.
<b>Management Group</b>	In the SAN/iQ software, a collection of one or more storage nodes which serves as the container within which you cluster storage nodes and create volumes for storage.



**Table 1    Glossary for the VSA**

<b>Term</b>	<b>Definition</b>
<b>Quorum</b>	A majority of managers required to be running and communicating with each other in order for the SAN/iQ software to function.
<b>Replication</b>	Synchronous replication of volumes across storage nodes in a cluster.
<b>Virtual Manager</b>	<p>A manager that is added to a management group but is not started on a storage node until it is needed to regain quorum. The virtual manager must be started manually.</p> <p>A virtual manager is added to a management group automatically when you create a management group with an even number of managers, for example, if using two storage nodes to create the management group.</p>
<b>Volume</b>	A logical entity that is made up of storage on one or more storage nodes. It can be used as raw data storage or it can be formatted with a file system and used by a host or file server.
<b>VSA</b>	The Virtual SAN Appliance for VMware® ESX Server, abbreviated as VSA.

## Designing Your Virtual SAN

The design of your virtual SAN affects both its performance and its reliability for production use. Use the concepts and practices in this manual when installing and managing your virtual SAN to ensure optimal performance and reliability.

## Configuration Requirements

The following prerequisites are required for LeftHand Networks to support your virtual SAN for production use.

- A qualified server running VMware ESX Server 3.0.x. Qualified servers are listed in VMware ESX Server 3.0.x systems compatibility guide, which can be found at [http://www.vmware.com/pdf/vi3\\_systems\\_guide.pdf](http://www.vmware.com/pdf/vi3_systems_guide.pdf).
- 1GB of memory reserved.
- A single virtual CPU with at least 2000 MHz reserved.
- A connection to a virtual switch with at least a Gigabit network connection.
- A virtual disk with 5 GB to 2 TB (2047 sectors) of space located on internal SCSI disk storage, or direct attached SCSI storage that is not accessible from more than one physical server. (Note that the SAN/iQ OS consumes a small amount of the available space.) This virtual disk must be connected to SCSI address 1:0.
- All virtual disks for the VSA must be configured as independent and persistent to prevent VM snapshots from affecting them.
- The VMFS datastore for the VSA must not be shared with any other VMs.
- Two or more VSAs on separate physical servers with 2-way replication, and a [Failover Manager](#) is the minimum configuration for high availability with automatic failover.
- Two or more VSAs on separate physical servers can be clustered with a [Virtual Manager](#) for manual failover capability.

## Best Practices

Other configuration recommendations are useful to improve the reliability and performance of your virtual SAN. Consider implementing as many of these best practices as possible in your virtual SAN environment.

Each VSA should

- Have a virtual switch comprised of dual Gigabit Ethernet or more. Providing network redundancy and greater bandwidth improves both performance and reliability.

- Be located on the same virtual switch as the VMkernel network used for iSCSI traffic. This allows for a portion of iSCSI IO to be served directly from the VSA to the iSCSI initiator without using a physical network.
  - Be on a virtual switch that is separate from the VMkernel network used for VMotion™. This prevents VMotion traffic and VSA IO traffic from interfering with each other and affecting performance.
  - Be configured to start automatically and first, and before any other virtual machines, when the ESX Server on which it resides is started. This ensures that the VSA is brought back online as soon as possible to automatically re-join its SAN/iQ cluster.
  - Use redundant RAID for the underlying storage of a VSA in each ESX Server to prevent single disk failures from causing VSA node failure. Do not use RAID 0.
- See the section “Planning the RAID Configuration” of the Storage chapter in the *LeftHand SAN User Manual* for detailed information about using RAID for individual server-level data protection.

The *LeftHand SAN User Manual* is available in PDF format and is installed in the Documentation subdirectory with the Centralized Management Console program files. You can also find the information in the online Help, available from the Help menu and the context-sensitive Help buttons on each window in the Centralized Management Console.

## Unsupported Configurations

Some configurations that are possible using VMware ESX and the VSA are specifically not supported for production use at this time.

LeftHand Networks **does not support:**

- More than 1 NIC configured on the VSA.
- Use of any shared storage as the data location of the VSA.  
This includes shared DAS, Fibre Channel, iSCSI, and NAS.
- Use of VMware snapshots, VMotion, HA, or DRS on the VSA itself.
- Use of any configuration that VMware does not support.
- Booting ESX off of a VSA cluster.
- Extending the data (SCSI 1:0) vmdk of the VSA.

- Co-location of a VSA and other virtual machines on the same physical platform without reservations for the VSA CPU and memory.
- Co-location of a VSA and other virtual machines on the same VMFS datastore.
- Use of VSAs on other VMware platforms such as VMware Server, Workstation, or Player.

## Hardware Design for VSA

The hardware platform used for a virtual SAN affects the capacity, performance, and reliability of that virtual SAN. The following hardware features specifically affect the VSA configuration:

- CPU
- Memory
- Controllers and Hard Disk Drives
- Network Adapters

### CPU

Because the CPU of the VSA must be reserved, platforms that will host a VSA and other VMs should be built with more processor cores to accommodate the additional VMs. Multi-core processors with at least 2GHz per core should be used so that a single core with at least 2GHz can be reserved for the VSA. All additional cores are then available for use with other VMs, thereby avoiding resource contention with the virtual SAN. For example, a platform with two dual core processors could host a VSA and use 3 cores to share for other VMs.

### Memory

Similarly the memory of the VSA must be reserved. For platforms that will host a VSA and other VMs, build in additional memory to accommodate the additional VMs. Assuming the hypervisor and management applications will use less than 1GB, all memory beyond 2GB is available to use with other VMs,

again avoiding resource contention with the virtual SAN. For example, a platform with 10GB of RAM could host a VSA and use 8GB of memory to share for other VMs.

## Controllers and Hard Drives

The internal disk controller and actual hard disk drives of a platform affect the capacity and IO performance of the VSA. Ideally VSAs should use storage that is hosted by many SAS or SCSI drives. If you are designing a new server that will host VSAs, you should

- Use as many hard drives as the platform will allow and prefer faster rotation speeds. The more hard drives and the faster their rotation speed, the more IOPS, and better performance.
- Select controllers with protected cache.

## Changing the Disk Size on the SAN

Use the VI Client if you want to change the disk size of a VSA that is currently in your LeftHand SAN.

You must follow these steps in the sequence listed or the VSA will be unable to participate in a management group.

- 1 Remove the VSA from the management group.
- 2 Using the VI Client, power off the VSA.
- 3 Select the VSA and select Edit Settings to open the Virtual Machine Properties window.
- 4 Select the Hard Disk that is on the port SCSI 1:0 and click Remove.
- 5 Click Add and use the Add Hardware Wizard to add a new hard disk of the desired size to port SCSI 1:0.

See detailed instructions in the section “Configure a data disk” in the *VSA Quick Start Guide*, installed in the Documentation directory with the Centralized Management Console files.

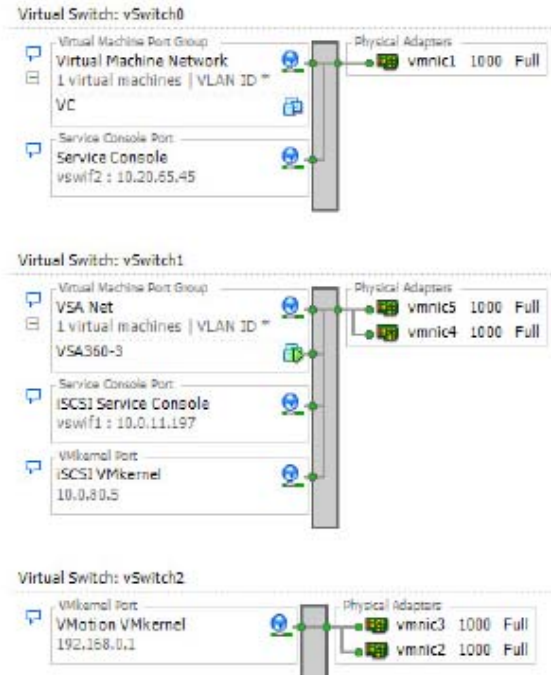
- 6 Power on the VSA.
- 7 In the Centralized Management Console, reconfigure RAID.
- 8 Add the VSA to a new or existing management group and cluster.

## Network Adapters

The number of network adapters available in a platform affects your options for configuring virtual switches. VSAs that will have a dedicated ESX Server platform only need 2 Gigabit network adapters. Platforms that will host VSAs and other virtual machines should have at least 4 Gigabit network adapters so that two adapters can be dedicated to the VSA and iSCSI traffic.

## Virtual Network Design

The virtual switch that is used for VSA should be at least a redundant Gigabit network, and should be shared with the iSCSI VMkernel network. If possible, an additional virtual switch with a VMkernel network for VMotion should be configured in order to separate VMotion traffic from iSCSI traffic. If enough network connections are available, the virtual network shown in [Figure 2](#) is an ideal example for a virtual SAN.



**Figure 2 Example virtual network for a virtual SAN**

Performance and reliability can be improved even further by using more than two Gigabit adapters in the iSCSI and VSA virtual networks.

## Using VSAs with LeftHand SANs

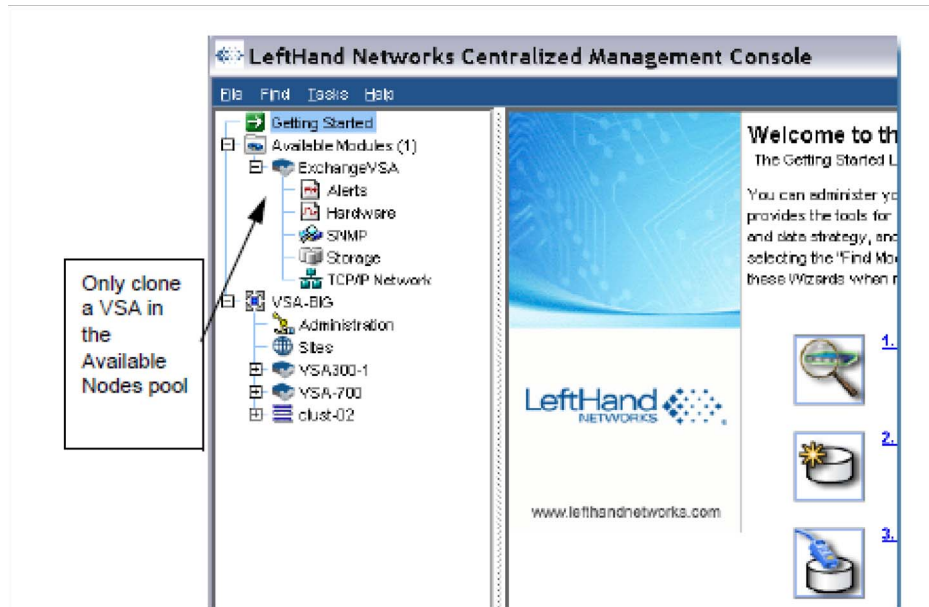
VSAs and physical platforms can be mixed together in SAN/iQ management groups and clusters. It is also possible to configure many VSAs with varying hardware architectures that will yield multiple performance points. When mixing virtual and physical platforms, and mixing differently configured VSAs, consider the following guidelines.

## Cloning VSAs

If you want to clone a VSA, you must do so while the VSA is still in the Available pool, before you add it to a management group. Cloning a VSA after it is in a management group is not supported.



**Tip:** Configure the first VSA for RAID, Alerts, SNMP and Networking. Then create your clone with the hardware settings already configured.



**Figure 3 VSAs in the Available Nodes pool**

## SAN/iQ Managers on Physical Platforms

When possible, all managers in a management group should be on physical platforms instead of VSAs, and ideally on the fastest platforms in the management group. While VSAs can run managers, physical platforms provide better performance and a lower likelihood that they will be rebooted for administrative reasons.



## VSA Managers on Physical Platforms

If VSAs are running managers, ensure that those VSAs reside on separate physical platforms. Otherwise, rebooting a single physical platform could cause a loss of quorum in the management group.

## Performance Impact of Mixing Platforms in Clusters

Mixing VSAs and physical platforms in the same SAN/iQ cluster is possible but yields unpredictable performance results. Mixing VSAs built from different hardware in the same cluster will also yield unpredictable performance. Typically the overall performance of the cluster is close to the aggregate average of all the storage nodes in the cluster.


Additionally, you may perceive performance degradation if a faster VSA fails, thereby causing volumes to fail over to a slower VSA. This type of performance degradation is normal in a mixed cluster.

## Sample Configurations

Some sample configurations and their uses and benefits are described below.

### Single Node Configuration

A single node configuration is useful for implementing a simple iSCSI target that can scale up to a redundant iSCSI SAN of any size in the future. A single node configuration consists of one VSA in a single management group, with a single cluster, and running a single manager, shown in [Figure 4](#).

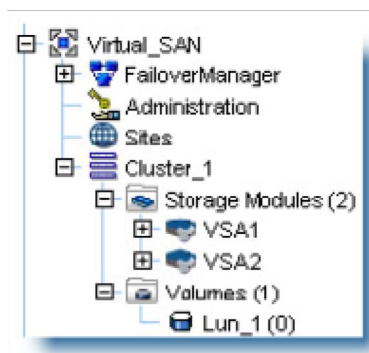


**Figure 4 Single node configuration**

## Two-Node Configuration

A two-node configuration is the smallest redundant configuration possible. Automatic failover between nodes requires a failover manager. A two-node configuration includes two VSAs in a single management group, with a single cluster, running two managers and with a failover manager added to the management group, shown in [Figure 5](#).


For more information about the failover manager, see [“Failover Manager Overview”](#) in the *Multi-Site SAN User Manual*, installed with the Centralized Management Console as a PDF in the Documentation folder.



**Figure 5 Two-node configuration**

## Three-Node-Plus Configuration

All configurations greater than two nodes can be redundant, running either 3 or 5 managers, and containing multiple clusters with as many VSAs as desired in each cluster, shown in [Figure 6](#).



**Figure 6 Three-node-plus configuration**

